

APPLICATION FOR UNITED STATES LETTERS PATENT

by

CHRISTOPHER R. STEPHENS

and

HENRI WELBROECK

for

**SYSTEM AND METHOD FOR CREATING, DISTRIBUTING
AND MANAGING ARTIFICIAL AGENTS**

CROWELL & MORING LLP
1001 Pennsylvania Avenue, N.W.
Washington, D.C. 20004
(202) 624-2500

Attorney Docket No.: ADT0001-US

1683323
Drawings: 1688949

SYSTEM AND METHOD FOR CREATING, DISTRIBUTING AND MANAGING ARTIFICIAL AGENTS

FIELD OF THE INVENTION

The present invention relates to intelligent artificial agents and more particularly to systems and methods for the development, distribution and management of intelligent artificial agents.

BACKGROUND OF THE INVENTION

Technological progress has allowed machines to replace humans in a growing range of repetitive tasks, improving the efficiency of industrial processes and raising productivity in a dramatic way. This improvement however has not been fully matched in the service sector, in part because many of the tasks carried out in the services industry are perceived to require intelligent reasoning and skills that are not easily emulated by machines.

This is particularly true in finance where fundamental analysis has been a human prerogative and traditional technical analysis has also relied on humans to interpret its results. There are many technical analysis tools available, but these are often relatively unsophisticated systems that are viewed as "toys" by more quantitatively minded analysts. While there are comparatively better systems, such as Metastock, available from Equis International, Salt Lake City, Utah, that provide a large amount of potentially useful information, these systems are not true prediction systems since they rely, ultimately, on heuristic rules that cannot have global validity over time in a constantly changing financial market.

A large part of finance is associated, for obvious reasons, with trying to predict what will happen in the future. Both fundamental and technical analysis are concerned with the question of prediction but approach it from completely different points of view. Traditional economic theory (the Efficient markets hypothesis) states that markets are by definition unpredictable. However, there is

mounting evidence that markets are not entirely efficient or rational; that is to say they are not completely random. Any system that can take advantage of such inefficiencies would be capable of permanent profit making. Unfortunately, inefficiencies are neither linear nor permanent in character. Additionally, the exploitation of an inefficiency can eventually destroy it. Hence, one is faced with a situation where a particular financial instrument exhibits only temporarily a “non-random” or “predictable” behavior.

Artificial Intelligence is a relative newcomer to the field of finance. However, there are now many systems available for financial prediction that contain elements such as neural networks and genetic algorithms. These are highly non-linear systems and use computationally complex algorithms whose results can be highly unstable. Additionally, they do not offer predictions a priori but need to be “tuned” or “trained” by the user who almost inevitably is not an expert in artificial intelligence and, therefore, likely to produce unreliable results. If the “tuning” or “training” is successful, the producer of the system can claim credit, and if unsuccessful the producer can blame the client for not training the system well enough. Additionally, training of neural networks or optimization in genetic algorithms, if done correctly, tend to be computationally intensive processes requiring computational resources and resources of time from the client that could better be dedicated to other tasks.

On the forefront of artificial intelligence research are intelligent artificial agent systems which are now opening new avenues for productivity increases in areas where humans are carrying out repetitive intelligent tasks. This sets the stage for a new technological revolution that will change the way in which many services are rendered. The spectacular growth of the Internet is one of the first indications that this process is already underway. Commercial applications of intelligent agents have essentially been restricted to “data mining” where a more

intelligent search of databases is carried out. In fact many such systems are no more sophisticated than standard web search engines.

Thus, there remains a need for a different approach to development, distribution and management (or use) of intelligent artificial agents, and particularly those that are applicable to dynamic systems such as stock markets where a crucial aspect of the agent system must be its ability to adapt to changing conditions and where prediction plays a fundamental role.

There are clear opportunities for using intelligent artificial agents as analytical tools, and particularly so in finance. In the Internet sector alone (i.e., online trading), as of 1999 there were more than 5.4 million active online brokerage accounts managing assets of nearly \$500 billion. This amount is expected to rise to \$3 trillion by 2003. Discount brokers are presently reaching the limit of how efficient they can make their operations and still make a profit in the midst of intense competition and falling transaction fees. Accordingly, there is an increasing tendency to add on new services that give substantial added value. In particular, investors continue to seek relevant and timely information that can give them a competitive edge in their trading decisions but at a price that does not excessively impact the overall cost of trading.

While the glory days of the full service broker will likely soon disappear, the need for advice on equity trading and other elements of financial assessment will not. Accordingly, there is a need for affordable advice in terms of both fundamental and technical analysis on equities and other financial instruments as well as a constant need for assessment of risk, tax positions, portfolio management, etc. Indeed, the potential market for timely and affordable advice is enormous. A conservative market penetration estimate of 1% of managed online assets by 2003 and total commissions and fees equivalent to 0.5% of assets managed generates revenues of \$150 million. Artificial agent systems are capable of offering superior

investment strategies at a cost substantially less than that of their human counterparts.

In addition, there are over 1000 large hedge funds operating in the United States alone managing \$260 billion in assets that could benefit from improved analytical tools.

Another market that would benefit from improved analytical tools is the so-called "market makers" arena. Market makers are extremely susceptible to adverse information. Traders that enjoy information that is not (yet) known by the market maker can take shares from the market maker and later resell them for a higher price when the information has become available. Individual clients that repeatedly "pick off" market makers in this way can be identified and traders at market making firms devote some effort to this identification task, but do so using only their subjective experience rather than a systematic tool. Indeed, a 10% increase in efficiency in detecting adverse information on an estimated 45% loss in revenue from the adverse information risk for a typical market maker with annual trading revenue of \$800 million leads to a savings of \$80 million. Therefore, there is always a need for improved identification procedures for market makers.

In view of the foregoing, there is a clear need for improved, timely and affordable analytic tools for both personal and business use.

SUMMARY OF THE INVENTION

The present invention provides systems and methods for the development, distribution, management and use of intelligent artificial agents.

In particular, the present invention provides a system and method for developing or creating, distributing and managing artificial agents that offer timely, relevant information to investors preferably by means of an agent leasing service.

In the currently preferred embodiment, the artificial agents are artificially intelligent, computer algorithms that provide superior trading strategies, which give added value to the investment decisions of any financial professional or individual investor.

More specifically, in accordance with the present invention the agent is an algorithm that executes a sequence of operations whose final result is a recommendation for trading a given financial instrument and an evaluation of the agent's expected future performance. The algorithm is created in an "agent factory" (discussed later herein) and is characterized by a set of parameter values. Preferably, these parameter values are written to a file which can then be transmitted to a user's computer via the Internet, or other suitable method of transmission, where a set of executable programs read these parameter values and recreate the agent and its associated trading strategy and recommendations.

The "agent factory" (discussed later herein) produces the files with the parameter values. Preferably, the agent algorithm (i.e., how the agent reads data from files and transforms this data into a sequence of operations) does not change when agents are replaced in light of the evolving market conditions.

In addition, in accordance with the present invention, a rental or leasing system is the preferred distribution system because there is no permanent solution to the problem of prediction. Thus, the present invention further provides a system and method for conducting an artificial agent leasing service.

Still further in accordance with the present invention, there is provided a simple and efficient mechanism whereby an individual or institutional investor is able to identify, select, order and enable particular artificial agents that are best suited for a particular investment strategy. That is, the present invention provides a means of managing a plurality of artificial agents.

More specifically, the present invention can be considered a consulting system that is based primarily on the decision-making abilities of artificial agents. The system includes a graphical user interface to interact with investors and aide in selecting the appropriate artificial agents. The artificial agents of the present invention replace, in many respects, human agents. Importantly, the artificial agents provide superior trading strategies, which are based on the detection of statistical biases in financial time series.

A constantly changing global marketplace demands prediction tools of the highest quality. Today's demanding financial professionals and investors need the power of improved technology. The present invention offers them this power with an artificial agent-based prediction system, wherein each agent is created from an adaptive system and is, in its turn, an adaptive system. More specifically, the present invention includes the artificial agent factory where agents are created. In the factory, agents exist in a competitive "Darwinian" environment where an ecology of agent strategies compete. Survival of the fittest in this context means survival of the most predictable. Each agent that leaves the factory includes a quantitative predictability value which is a direct measure of the agent's "fitness" and which can be used by a user to help choose which particular agent or agents to employ.

With respect to artificial agent distribution, the leasing system of the present invention makes it possible to provide pre-trained, computationally powerful prediction agents adapted to current market conditions without placing an excessive burden on the user's own resources. Through this innovative solution years of intensive research can be made available to financial decision makers worldwide. Preferably, the artificial agents provide advice on buying and selling financial instruments such as equities.

The artificial agents in accordance with the present invention are preferably trained to recognize emergent price structures (i.e., "predictability bubbles"), design

profitable trading strategies to exploit them and anticipate their end by withdrawing before profits turn to losses. Constant agent renewal, in accordance with the present invention, results in improved overall predictability even in view of rapidly changing market conditions.

In accordance with a preferred embodiment of the present invention, the life span of any individual artificial agent is limited. That is, the strategies on which a particular artificial agent is based will eventually be retired and be replaced by new strategies. This is a crucial aspect of the present invention as permanent, static solutions to evolving problems do not exist. This is especially true in financial markets where a useful strategy at one moment in time may become obsolete at another due to changing market conditions and where an overly successful strategy carries the seeds of its own destruction due to excessive market impact. Artificial agent production is therefore a continuous adaptive process that is computationally intensive and preferably carried out at a central computing and research facility, i.e., the agent factory.

With the artificial agent leasing service of the present invention, any application based on the agents' recommendations will likely depend on the continuing stream of fresh adapted agents. Thus, the agent factory of the present invention preferably supplies the appropriate number of agents at the appropriate time.

In accordance with the currently preferred embodiment, the present invention preferably provides a variety of agent-based tools for financial decision making including:

- an investor tool for individual and institutional traders;
- a fund tool for fund managers;
- a market maker tool for market makers; and

- an agent editor tool to allow a user to create an agent with the characteristics of his choice.

Within the finance arena, similarly-styled intelligent, agent-based techniques are applicable to other output variables (interest rate derivatives, currency derivatives, etc.) and other inputs including fundamental analysis data, news clippings, etc. The more this type of information is quantifiable, the more it can be used to augment the decision-making skills of the artificial agents of the present invention.

There is, of course, a need for improved artificial agent technologies in other areas. Accordingly, the agent-based tools of the present invention are also applicable to disciplines as diverse as the electricity supply industry and the insurance industry, where a search for and exploitation of predictability bubbles may be implemented.

Agent based modeling has drawn increased attention in recent years, in part for its value in helping to understand the behavior of complex adaptive systems, such as financial markets, and learn to forecast their evolution. For example, in portfolio management or day trading, artificial agents offer superior trading strategies and analysis when compared to their human counterparts.

The artificial agents of the present invention preferably include intelligence derived from a predictability value which itself is preferably further based on neural network information.

With respect to predictability, it is legion that “past performance is no guarantee of future returns.” However, a predictability indicator provided by the present invention has an objective criterion that links the past and future and directly measures an agent’s expectation of future success. The predictability

function provides direct expectation of gains rather than anecdotal evidence based on past performance.

With respect to neural networks, the present invention preferably augments each artificial agent with the results of a self-organized memory-recalling neural network that provides to each agent the capability to learn via “avalanches” of neural activity.

Investor Tool

The investor tool is an agent based system offering trading advice and superior trading strategies to investors who trade equities (or other financial instruments). Agents are equipped with functionality to predict price movements of equities (or other financial instruments). As every instrument has its own unique pattern of bubbles of predictability, agents are custom made and dedicated exclusively to the trading of a particular instrument. As noted above, the agents are created by a central agent factory and thereafter preferably leased to clients (or users) for a monthly fee, for example. Due to continuous changes in market conditions, agents have a finite lifetime whereupon they are retired and returned to the agent factory for retraining to take into account the new, prevailing market conditions. In this way, agents, and the trading strategies they employ, are better able to adapt to changing market conditions.

In the agent factory a systematic search for valid technical signals (templates) and their associated neural interpreters is carried out. The products of the factory are deliverable prediction agents, which are capable of transforming transitory predictability bubbles into sharp yield curves. The user can either follow a single agent's yield curve by systematically matching his trading decisions, or compare the opinions of several agents. Management software allows the user to operate and consult agents and view their decisions and historical yield curves, as well as other information.

The investor tool is well suited to institutional investors or individual investors. With individual investors, the present invention is preferably implemented on the Internet. Investors in the now extremely large "on-line trading" community actively seek high quality, timely information, but at a cost that does not excessively impact the overall low cost of trading. The present invention satisfies the simultaneous requirements of high quality, timely information, instantaneous 24 hour access and low cost that cannot be profitably offered by conventional human agents. In a preferred embodiment, users consult with artificial agents on-line as if they were asking advice from a human analyst.

Thus, the on-line implemented investor tool is directed to individual investors and provides information useful to analyze possible future risks and benefits as part of the process a private investor goes through in confirming his intention to trade.

Fund Tool

This tool preferably supports automatic Fund operations and is, essentially, a generalization of the investor tool described above. Instead of focusing on individual stocks as in the investor tool, the Fund tool preferably provides a complete trading strategy, including the use of derivatives if required, and interfaces with a user's order routing and execution systems.

In this embodiment of the invention, the agents' trading recommendations preferably are accessed by a separate functional unit that also loads in the current state of the fund's portfolio and outputs suggested trades that would improve this portfolio's future performance or reduce risk. This separate functional unit (an artificial agent) preferably has the following functions:

- a) seeks a consensus among the various agents for each individual instrument;

- b) decides how much money should be assigned to this consensus decision in light of the fund's present position in this instrument and in other instruments; and
- c) approves or rejects the final trade according to user-specified criteria relating to the fund's guidelines such as maximum exposure allowed per instrument type, or other known criteria.

The fund manager can decide to authorize this agent for automatic execution, in which case the recommended final trade is automatically passed on to an order execution agent. The order execution agent submits the order to an electronic trading system, manages the dialogue with this trading system and reports when the order is pending match or done. Reporting involves a message to the user (in this case a manager) and an automatic update of the user's portfolio database to reflect the new position.

Market Maker Tool

This tool seeks to confront the "adverse information cost" incurred by market makers when they do not immediately become aware of new information and respond to it by moving their quotes. This information can be technical in nature (e.g., price patterns that are deemed to be relevant by day-traders), news clippings, movements in the market makers montage, or information on trading decisions by large institutions, for example. Researchers estimate that about half of the quoted spread is lost to the adverse information risk. Much of this adverse information is in fact indirectly available to the market maker, to the extent that one can filter it out of the vast amount of incoming information describing everything from the market maker montage to customer order flow. This filtering task is similar in essence to the task of screening agents in the agent factory for the investor tool.

The stream of incoming orders at a given market maker provides valuable information which can be used to anticipate short-term price fluctuations and warn

the market maker to cover his positions. The present invention is ideally suited to the task of filtering out relevant information from this order flow by identifying individual clients or client groups that are consistently picking off the market maker, in the same way that the investor tool technology identifies agents that are likely to generate future profits.

Agent Editor Tool

The agent editor is a system and method allowing a user at a remote terminal to create agents that implement a trading strategy desired by that particular user. That is, this embodiment of the present invention permits a user to tailor an agent to a particular trading strategy which might be tied to an individual investor's trading philosophy.

It is therefore an object of the present invention to provide a system and method for creating, distributing and managing intelligent artificial agents.

It is another object of the present invention to provide artificial agents that are trained to aid decision making by investors in financial markets.

It is yet another object of the present invention to provide an agent factory that continually generates new artificial agents each having a quantitative predictability value.

It is still another object of the present invention to provide an artificial agent leasing service.

It is another object of the present invention to provide a management system for a user of the artificial agents.

It is another object of the present invention to make artificial agents available to users over the Internet.

It is another object of the present invention to provide artificial agents suitable for individual and institutional investors, hedge funds and market makers.

It is another object of the present invention to provide an agent editor tool that permits a user to customize an agent's trading strategy.

It is still another object of the present invention to provide for each of a plurality of artificial agents a predictability value based on a reconstructed predictability landscape.

It is another object of the present invention to provide self-monitoring artificial agents that are capable of retiring themselves when it is determined that their trading strategy is no longer effective.

It is yet another object of the present invention to provide artificial agents based on different analysis templates.

It is still another object of the present invention to provide a plurality of artificial agents to users and further provide a recommendation as to which agents of the provided agents to employ.

These and other objects will become apparent to those skilled in the art upon reading the detailed description of the invention in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a general schematic diagram of a preferred embodiment of the present invention.

Figure 2 is a flowchart illustrating the identification of variables that are relevant to landscape reconstruction in accordance with the present invention.

Figure 3 is a simplified three-dimensional depiction of a landscape reconstruction.

Figure 4 is an exemplary curriculum vitae (CV) for an artificial agent in accordance with the present invention.

Figure 5 illustrates a stock manager user interface screen in accordance with the present invention.

Figure 6 illustrates an agent manager user interface screen in accordance with the present invention.

Figure 7 illustrates a ticker user interface screen in accordance with the present invention.

Figure 8 illustrates an agent round table user interface screen in accordance with the present invention.

Figure 9 illustrates the stock manager user interface screen including an agent recommendation window in accordance with the present invention.

Figure 10 illustrates graphical analysis user interface screens in accordance with the present invention.

Figure 11 is a depiction of an agent editor tool user interface screen in accordance with the present invention.

Figure 12 is an example of the possible scenarios produced by the agent editor tool in accordance with the present invention.

Figure 13 is a flowchart illustrating how it is determined that an artificial agent should be retired for retraining in accordance with the present invention.

DETAILED DESCRIPTION OF THE INVENTION

The present invention relates to intelligent artificial agents and more particularly to systems and methods for the development, distribution and management of intelligent artificial agents.

The present invention is designed for systems that are in a constant state of flux, or, more particularly, are adaptive. Accordingly, the present invention is particularly suited to financial markets, which respond to the actions of various market participants. It has been found that, with respect to stocks for example, structures continuously emerge in price charts as diverse trading strategies interact to produce coherent movements. Representing the combined effects of all these trading activities, these structures are an example of a phenomenon known in physics as “emergence of collective excitations.” In finance, researchers refer to emerging structures as “predictability bubbles.”

Even the best financial traders eventually lose step when markets are adapting quickly. However, an intelligent artificial agent in accordance with the present invention can recognize predictability bubbles when they arise, track their evolution and, significantly, anticipate their end. Thus, the intelligent artificial agents (hereinafter “artificial agents” or “agents”) of the present invention operate to identify price predictability patterns in equity markets and provide specific recommendations to a user as to how best to exploit that pattern. Each agent preferably is self-monitoring in that the agent tracks the evolution of a predictability bubble and anticipates its (the agent’s) own future performance.

Figure 1 shows a general schematic diagram of a preferred embodiment of the present invention, which, in this case, is directed specifically to a financial services implementation. However, it should be evident to those skilled in the art of artificial agents that the artificial agent development, distribution and management methods and systems discussed herein are equally applicable to other disciplines

where decision making based on a plurality of variables might be enhanced by the use of artificial agents, e.g., medical expert systems, the electricity industry and the insurance industry.

Referring now to Figure 1, which provides a broad overview of the present invention, artificial agents 110 are created in an agent factory 115. As will be explained in more detail later herein, agent factory 115 is a complex adaptive system that is capable of creating artificial agents 110 on an as-needed basis. Specifically, agent factory 115 analyzes both past and current market data 120 (e.g., stock market data) and creates new agents 110 when it is determined that previously created agents 110 no longer perform in a predictable manner.

As new agents 110 are created by agent factory 115, they are distributed to a subscribing client or user, and preferably to a management system 125 that manages a plurality of artificial agents 110 for the user. Management system 125 preferably includes a user interface through which the user may identify, select, order and put into operation individual artificial agents 110, as desired. The user interface preferably also provides the user a means of monitoring the performance of any individual agent over its lifetime. In practice, management system 125 preferably is also in contact with market data 120 such that the performance monitoring function is a real time function and real-time advice or recommendations are effectively provided. Agent factory 115 preferably also comprises management system 125, or at least the functionality thereof which is described in detail later herein.

Also, as shown in Figure 1, the present invention provides, with respect to each implemented artificial agent 110, a specific recommendation 130, which in the financial industry context, might include a recommendation to buy, sell, or stay neutral in a position with respect to a particular stock listed on a particular stock exchange. Figure 1 also shows an automated trade clearing system 135, which preferably is in communication with management system 125 to thereby receive the

various recommendations (shown via dotted line in Figure 1) and automatically execute the recommended action. When an automated trade clearing system is not implemented, a user preferably receives the aforementioned recommendation and uses it exclusively, or as part of the basis for his ultimate trading decision making.

Self-monitoring artificial agent

A self-monitoring agent in accordance with the present invention is an algorithm that preferably takes as input a time series, such as a price series for a given financial instrument, and executes a sequence of operations whose final result is a (set of) recommendation(s) on appropriate actions to be taken by the user and a measure of the expected quality of these recommendations, i.e., the expected future performance of the agent. In the preferred embodiment of the present invention the output is a trading strategy that produces recommendations for trading that instrument, and the quality of the recommendations is equal to the profits that the agent expects will be earned by applying this trading strategy. In this case the future performance is directly related to the degree of predictability of the time series being considered and the future performance estimate is called “predictability” (see below for a discussion of the latter). The agent algorithm is created in the agent factory (discussed below) and preferably comprises two distinct algorithms: an A-brain and a B-brain. The former searches through a large library of technical templates looking for one of high predictability, while the latter relies on a neural network to modify, when necessary, the result of the A-brain.

The technical template preferably is a technical trading strategy itself such as a moving average, MACD, etc., which are familiar to those skilled in the art. This A-brain strategy gives rise to a yield curve which is then fed to the B-brain which comes from a library of neural networks with, for example, a self-organized critical learning algorithm or a conventional backpropagation algorithm (many of which are commercially available) and well-known to those skilled in the art.

Preferably, the purpose of the A-brain is to generate yield curves that have desirable properties, such as low volatility and high yield per unit time. The purpose of the B-brain is to detect changes in market tendency that would invalidate the A-brain strategy, i.e., to identify when a predictability bubble associated with a given strategy is about to end, and to amend the A-brain strategy if and when necessary. Thus, the A-brain strategy might yield a buy signal but be overruled by the B-brain which may negate or overturn the buy signal. The A-brain and B-brain together comprise an agent and yield a trading strategy and a subsequent yield curve that may be described in terms of a set of parameters including, for example, timescale(s) of the technical template used, predictability, aggressiveness, the duration in time units of the agent's strategy (i.e., the duration of its past trading), the agent's precision (i.e., the percentage of correct decisions in past trading), the total number of decisions taken in past trading, the total profit of the agent in past trading, force threshold (see below), type of B-brain used, etc. Also relevant are several "geometrical" variables associated with an agent's "trading canal". The latter is defined by taking the agent's yield curve for its past trading and finding a pair of parallel lines that give a best envelope to this curve (see for example the Value curve in Figure 5). Parameters preferably include the length, slope, initial and final width of the canal. Preferably, all or a subset of these parameter values encode the agent and, therefore, its complete trading strategy such that when the agent is transmitted via internet or other means of communication to a user, the values act as input parameters to a set of executable programs which recreate the agent and therefore trading strategy and recommendations associated with those parameters.

Predictability and Landscape Reconstruction

In accordance with the present invention, each agent is preferably self-monitoring such that it is capable of anticipating the evolution of its own

performance by analyzing the strength of the predictability bubble on which its strategies are based. To implement this functionality, the present invention includes a method for measuring predictability. Predictability is based on mutual-information-based reconstruction of a multivariate fitness landscape which evaluates an agent's ability to generalize into the future. This landscape also plays an important role in the evolution dynamics of agents at agent factory 115, as will be explained in more detail later herein.

In the context of financial markets, landscape reconstruction is the process of analyzing past market data and determining a predictability value for each agent or technical template (see below for more discussion of technical templates). More specifically, the purpose of this landscape reconstruction is to establish the correlation between an agent's past performance and its expected future performance. "Predictability" is by definition the statistical expected value of a target variable that measures future performance, such as the agent's profits over the next 10 days, the agent's profits until its next trade or some other performance measure. The predictability of a technical template is measured by creating an agent that applies a simple technical trading rule, such as "buy when the price is above the moving average", and measuring the predictability of that agent. To reconstruct the predictability landscape, one performs a large number of experimental agent runs using historical data. In each agent run one executes the following steps:

- 1) Choose an agent or create one from a template-based technical trading rule, e.g. 10 day moving average, Fourier with period 30 days, MACD, or any other of a number of well-known technical trading rules.
- 2) Compute the earnings (yield) curve that this agent would have produced had he followed this trading strategy over a particular trading window (e.g. data lines numbered 400 to 500 in an historical database where each data line may represent a day, an hour or a market tick).

- 3) Compute the value of the target variable (i.e., earnings, in this case) by continuing to apply the same strategy in a subsequent testing window (e.g., data lines numbered 501 to 510).
- 4) Write to a database:
- (a) The characteristics that represent the class of agent strategy being applied, i.e., the technical trading rule that the agent is using. Preferably, a distinction is made between templates that seek trends and those that seek an oversold or overbought position.
 - (b) The value of various variables describing the stock's price history, such as volatility, price range over the trading window divided by current price, or other variables that may have a correlation to a stock's price.
 - (c) The value of variables that represent the agent's yield curve, such as the agent's profits in the last 10 days, the duration in time units of the agent's strategy, i.e., the duration of its past trading, the agent's precision, i.e., the percentage of correct decisions in past trading, the total number of decisions taken in past trading, the total profit of the agent in past trading, the expected profit of the agent assuming it makes profits at the same rate as in the past, a linear regression coefficient for the yield curve, etc. Also preferably written to the database are several "geometrical" variables associated with an agent's "trading canal". The latter is defined by taking the agent's yield curve for its past trading and finding a pair of parallel lines that give a best envelope to this curve (see for example the Value curve in Figure 5). Some of the parameters are the length, slope, initial and final width of the canal as well as the behavior of the yield curve within the canal, e.g., how many times

the value curve approaches or touches the canal lines, or the profit since the last contact between the value curve and the trading canal.

- (d) The value of the output variable will normally be profit in the testing window. The latter may be represented variously depending on the agent characteristics one is trying to optimize. Preferably, it represents a fixed time interval such as 5, 10, 15, etc. days. It may also, for example, represent earnings until the agent dies or until the agent makes a certain loss (stop/loss mechanism) or any combination of the aforesaid mechanisms.

5) Repeat steps 2-4 for different windows, e.g., lines 410-510 etc.

6) Repeat steps 1-5 for different agents having different technical rules.

The outcome of steps 1-6 is a large database, D, of experimental runs giving, for each run, the input variables discussed in 4) above representing the information on agent type, price and past earnings curve, and the output variable. The landscape reconstruction process inputs information from this database and performs the following steps as illustrated in Figure 2:

- 1) Choose initial input variables such as discussed above.
- 2) Select screening criteria on the variables such that each has a more limited domain. In other words, each of the variables in 4) above has a characteristic initial domain. For example, the range of trading canal length (past trading) may vary between 10 and 300 time units. However, to optimize the output variable it is often desirable to restrict the domain of each variable to a restricted sub-domain. For example, it may be useful to restrict past trading to the domain 30-250 to ensure adequate statistics (i.e., there may be few available data points in the regions 10-30 and/or

250-300, or to eliminate particularly bad sub-domains that lead to negative values of the target variable).

- 3) Identify relevant variables. In this preferred process, one identifies the subset of variables in 4) that correlate strongly and robustly with the output variable over a given time period. For instance, one may find that the output variable is approximately linearly related to a particular input variable over many different testing windows or alternatively for another variable that there are no significant correlations over time. The former would be considered a relevant variable and the latter an irrelevant one.
- 4) Based on the space S formed by the set of relevant variables found in 3) over the restricted domains found in 2) one performs a landscape reconstruction by dividing the space S into elementary cells and assigning a predictability to each cell as the average of the predictabilities associated with the subset of D that corresponds to that cell. Once the predictability landscape has been fixed, any new template can be assigned a predictability by determining in which cell the template fits and assigning the associated predictability. The resulting predictability is preferably a value that varies between 0 and 2 for any strategy that can later be used by a live agent.

A simple landscape reconstruction with only two input variables (earnings in past 10 days and width of trading canal) is illustrated in Figure 3.

Agent Factory and Agent Characteristics

In accordance with the present invention, each agent 110 is created within agent factory 115. Since agents 110 have limited life spans due to ever changing market conditions, agent factory 115 preferably continually generates new agents having predictability values that are above a minimum level. In a preferred

embodiment of the invention, as noted previously, predictability ranges from a value of 0 to 2, with 2 being the highest predictability.

Agent factory 115 comprises a plurality of computer algorithms that perform the various functions described below. Those skilled in the art will recognize that two or more of the following functions may be combined in a single subroutine or procedure of the computer algorithm. Thus, what is important is the functionality described below and not necessarily how it is precisely implemented.

In accordance with the present invention, each artificial agent 110 is a prediction algorithm designed to recognize and exploit predictability bubbles. To create agents agent factory 115 of the present invention includes a plurality of software routines that together train the agents and determine which agents should be provided to users.

In a preferred embodiment of the present invention, agent factory 115 preferably comprises 12 functional software routines:

- 1) data,
- 2) run agents,
- 3) refresh,
- 4) trades,
- 5) portfolio,
- 6) test,
- 7) quality,
- 8) distribute files to user database,
- 9) temp,
- 10)select,
- 11)A-brain,
- 12)B-brain,

The function of each of these routines is discussed below.

Agent factory 115 preferably maintains two pools of agents: active agents that are made available to users and reserve agents that preferably are used to replace active ones when they are retired as a result of changing market conditions. The main functions fall in two groups: management of active or live agents, and production of new reserve agents.

First group: managing the pool of live agents

“Data” reads new daily market price data from a financial data provider (market data 120) and makes this information available to the other factory functions.

“Run agents” updates each existing agent’s recommendations in light of the new data and retires agents that have a negative predictability value.

“Refresh” looks for agents that have been marked as “retired” due to changing market conditions, places these in a “cemetery” database and replaces them with reserve agents (previously generated) that have a high predictability value and trading strategies that are diverse to those currently adopted by existing active agents, usually because they are using different prediction horizons or templates.

“Trades” looks for agent consensus decisions using various non-linear consensus rules that identify teams of two or more agents that agree to recommend a trade. For example, if there exists for a given stock an agent with a short term horizon and positive PFI (a so-called “market timing” agent) together with a medium to long term horizon agent and both with predictability above a certain threshold and there exists no similar agent with a contrary opinion, then a position will be opened if both agents are in agreement. Trades also reviews trading positions recommended by teams of agents in previous runs of the factory process and decides if such positions should be closed based on the new data.

“Portfolio” looks for agent consensus decisions based on a formula to arrive at a consensus decision taking into account the opinions of all agents. For example, this consensus can be derived from a risk and reward analysis. In this implementation, a consensus force is the weighted average of the positions recommended by each agent in that stock (e.g., 1 for a buy recommendation, 0 for neutral and -1 for a sell recommendation), where the weights are determined by an iterative algorithm that seeks to maximize the ratio of the portfolio’s yield to its volatility. “Portfolio” manages a set of virtual portfolios that calculate their positions in each stock from the consensus force signal, for example using a proportionality rule, and apply different investor strategies, including, preferably, both investor portfolios and leveraged derivative portfolios.

“Test” measures performance statistics by combining data from the cemetery database, an active agents database, a Portfolio database and a database of consensus trade decisions, such as the average profit earned by each agent in the past month and the average profit per share traded using different consensus trading rules, or other similar comparison paradigms.

“Quality” checks that each parameter of every active and reserve agent has the appropriate format and is within lower and upper bounds that ensure that they will perform according to a user’s expectations.

Second group: creating new reserve agents

“Temp” is a template analyzer that tests the effectiveness of different technical trading rules on a window of historical data and evaluates the predictability of each trading rule.

“Select” looks for predictable trading rules from those discovered by temp and selects a subset of these that are most predictable and satisfy diversity enforcement criteria that, for example, avoid repeating rules that are already being used by previously-created agents.

“A-brain” preferably creates a preliminary artificial agent that uses the technical trading rule selected and generates a preprocessed signal (chart) available for further processing.

“B-brain” preferably iterates a process that implements a pre-trained neural network from a library of candidates and completes the agent’s training by letting it learn how to interpret the preprocessed signal generated by “A-brain”; for each choice of neural network it creates an agent and analyzes its predictability. The most predictable agent from these trials is kept in a database to be later “packaged” and sent to the ultimate user. Among the B-brains in this library are some candidates that apply strategies that would be recognized as common technical trading rules by those skilled in the art, and preferably other candidates that apply more complex interpretations of the pre-processed signals from the A-brain, including contextual learning and backpropagation neural networks that are also known to those skilled in the art.

In a preferred embodiment of the invention, each agent is provided with a Curriculum Vitae (CV), as shown in Figure 4, with which the user may evaluate the characteristics of the agents, judge its performance and select the agent most suited to his needs. Included are basic data such as the agent name, which technical template it uses, its horizon, its neural network training algorithm, the date on which the agent graduated from the agent factory and the agent’s aggressiveness, i.e., how often, on average, it recommends a trade every 100 days, or other fixed period of time.

The second part of the CV relates to the agent’s performance during its training in agent factory 115 and consists of the agent’s precision, i.e., of all the recommendations it made what percentage was correct, the actual number of decisions taken and the actual number of correct decisions. Correct in this context signifies that the price rose when the agent recommended to buy or the price fell when it recommended to sell. Also preferably provided is the average percentage

profit per 10 days the agent is making with its strategy and finally the agent's predictability value when it graduated from the factory.

The third part of the CV shown in Figure 4 preferably indicates the same statistical measures as the second part, but now taking into account not only the performance during training but also performance during the time with the user. These numbers change according to whether the agent is performing better or worse than its average during the training phase. Also present are the agent's current recommendation and PFI (discussed in detail later herein).

The CV also states the number of similar patterns used by the neural learning algorithm of the agent associated with the present price data. When the number of similar patterns used is relatively large, this information lets the user know that the agent is adapting to a change in price structures and may modify its trading strategy in the future.

If an agent is retired, preferably displayed in the "comments" section is the reason for the agent's retirement. Agents can be retired for at least three reasons:

1. Low predictability: the agent detects that his expected future performance will be low.
2. High risk: the agent has made profits faster than anticipated and foresees the possibility of a (possibly temporary) correction.
3. Excessive neural activity: agents can detect a change in the environment (market conditions) and revise a large amount of relevant historical data in an effort to modify their trading strategies to adapt to the new situation. If this adaptation effort leads to an excessive change in the trading strategy, it is considered necessary to observe how this new trading strategy performs before the agent is used.

Agent Management and Utilization

In accordance with a preferred embodiment of the present invention there is provided a means to identify, select, order and manage or keep track of the various artificial agents 110 that are provided by agent factory 115. Such a system is shown generally as element 125 in Figure 1. The following discussion is directed specifically to the investor tool (either individual or institutional), but it should be appreciated by those skilled in the art that analogs of this scheme could easily be applied to both the fund and market maker embodiments of the present invention as described above.

Agent management in accordance with the present invention is preferably implemented in an electronic computer. The computer preferably is capable of storing a plurality of artificial agents 110 and, in conjunction with real-time market data 120, notifying a user of a recommendation of one or a plurality of artificial agents 110.

In a preferred embodiment of the invention, agent management system 125 accesses market prices via an interface to established data vendors (e.g., Reuters) using methods known to those skilled in the art, and preferably "calls" the agents at regular intervals of time, using an updating frequency that can be chosen by the user, to update the agents' recommendations and predictability values. Preferably, agents with a low predictability value are automatically retired. Further, management system 125 preferably call the agents at the end of a trading day with a request to review the daily events and apply neural network learning algorithms. In addition to the above functionality, which is preferably accomplished without a user's intervention, agent management and utilization is further accomplished via a plurality of user interface screens as illustrated by Figures 5-10.

a stock manager (Figure 5);
an agent manager (Figure 6);

a ticker (Figure 7);
an agent round table (Figure 8);
an agent recommendation (Figure 9); and
graphical analysis (Figure 10).

As will be appreciated by those skilled in the art, some of the data depicted on one screen might also appear in other screens. This facilitates the use of management system 125 and provides reassurance to a user that the underlying data is consistent. It is noted that Figures 5-10 are based on selected stocks listed on the Mexican stock exchange and artificial agents 110 for those stocks.

Figure 5 depicts a stock manager user interface screen in which stocks for which agents are available are listed under the name column. When a particular stock is selected, then, as shown in the lower left of Figure 7, all of the agents for that selected stock (or instrument) are listed. In a preferred embodiment of the present invention, agent factory 115 supplies a user with no more than five artificial agents 110 at any given time and for any given instrument.

For each agent listed in Figure 5, listed also is the predictability value for the agent, its state, i.e., is the agent alive (normal) or retired, and what its profit per 10 days is. As shown, there are preferably also buttons (operable via a mouse, for example) that permit a user to select individual or all agents and to request the agents' respective recommendations. Finally, an agent can be shown graphically. These latter two features are explained in more detail later herein.

Still referring to Figure 5, as already explained, the predictability value preferably is a number that varies between 0-2, 2 being the maximum predictability. Any number between 1-2 generally signifies a relatively good agent. A predictability value that is less than 0 signifies an agent that has been retired and preferably should no longer be used. Of course, the range 0-2 could be modified to any standard range, e.g., 0-10, that might be desired.

Referring still to Figure 5, the stock manager user interface screen also includes data such as the liquidity of the stock, industry sector of the stock, and other well-known parameters used to describe and/or analyze a particular stock such as earnings per annum (E12M), book value (BV), quarter when the last report was emitted (L. Rep.), price-earnings ratio (P/E), price to book value (P/BV), current price, previous price, current volume and book value to earnings (BV/E). Accordingly, this user interface screen allows a user to select a particular stock according to given fundamental/technical criteria and identify the agents associated with that stock.

The agent manager user interface screen of the present invention is depicted in Figure 6, and is similar to the stock manager, but is arranged primarily by agent rather than instrument. Accordingly, agents can more easily be selected. For instance, an agent can be listed by order of graduation date ("G Date" in Figure 6) whereby only the most recently issued agents might be selected. The agent manager user interface preferably lists all possible agents in the system, i.e. the number of instruments x 5 (assuming 5 agents are available for each instrument). In addition, the agent manager user interface screen preferably includes columns entitled price feedback indicator (PFI) and horizon (preferably listed in weeks), which indicates when the agent is expecting to find a profit-taking opportunity.

PFI is determined by presenting an agent with different hypotheses about the price during the next trading period. Each agent preferably is presented with 50 hypothetical prices regularly spaced on a logarithmic scale ranging from a 50% drop to a 50% rise from the last price. For each hypothetical future price the agent's recommendation is calculated. Of special interest here is the change between the actual recommendation at the last price and the hypothetical recommendation that the same agent would give in each future price scenario: for example an agent that is recommending to "buy" at the current price may decide to recommend to "sell" if the price were to rise 10% or more.

In general, the agent's first change in position on a rising price can be either "buy" (B), "sell" (S) or "no change" (N). Similar notation is used for a falling price. The table below indicates how a value is assigned to PFI in each case:

	Rising price	Falling price	PFI	Explanation
i)	B	B	>	support
ii)	B	N	+	short, stop loss
iii)	B	S	+	unstable
iv)	N	B	-	short, take profit
v)	N	N	n	will hold
vi)	N	S	+	long, stop loss
vii)	S	B	-	stable
viii)	S	N	-	long, take profit
ix)	S	S	<	resistance

where, i) indicates the existence of a support line. In other words even if price falls it is likely to rebound from said support line; ii) represents a stop loss on a short position, i.e. at some point if price is rising the agent will buy to exit from the short position. iii) represents an agent that is currently holding a neutral position and will buy if price rises but sell if it falls. iv) represents an opportunity for profit taking in a short position, i.e. on falling price one exits the short position with a buy and takes profits. vi) is the opposite of ii) representing a stop loss on a long position. viii) is the opposite of iv) representing an opportunity for profit taking by exiting a long position. ix) represents the opposite of i) indicating the existence of a

resistance line. That is to say even if price rises, it is likely to reach no higher than said resistance line.

Further, listed on the agent manager user interface screen shown in Figure 6 is the state of the agent 110 and its aggressiveness. Aggressiveness, in accordance with the present invention, preferably is a number that varies between 0-100 and indicates the number of trades every hundred days, on average, that the agent recommends. Accordingly, an aggressiveness value of between 25-50 is considered relatively aggressive. In contrast, a value under 10 is considered to be a relatively conservative agent.

The next column in Figure 6 is the template on which the agent is based. As explained earlier, an agent 110 is composed two brains: an A brain and a B brain. The A brain is based on one of a plurality of technical templates that seek structure in the price series. The templates include: buy and hold, moving averages, MACDs, momentum, stochastics, oscillators, composite moving averages, fourier analysis, candlesticks and the like, all with different characteristic time scales. It is noted that some of the templates look for tendencies in a price series (generically linear trends) and some look for oversold or overbought positions. For example, the moving average could be 3 day, 5 day, 10 day, 30 day, etc. The above templates are well known strategies to any practitioner in technical analysis. However, in contrast to conventional technical analysis, the present invention provides at least two important differences. i) As opposed to strategies based on fixed heuristics, e.g., buy when price crosses the moving average from below, the present invention can take the contrary strategy, i.e., sell when price crosses the moving average from below. Thus, for every conventional strategy there is an inverse strategy considered by the agents of the present invention, and ii) strategies preferably are evaluated strictly according to their predictability.

The important point, as discussed previously, is that the template that maximizes the predictability function is chosen at agent factory 115 and passed to

the user. As explained above, the B brains portion of an agent preferably is a neural network. The configuration of the neural network is also listed on the agent manager user interface screen shown as "Forest Fire" and "Quicksand" in Figure 6 and in the agent CV of Figure 4. Finally, the agent manager user interface screen includes the predictability value for each agent and its profit per 10 days which indicates how much the agent has been earning on average every ten days.

Preferably, as shown in Figure 6, this particular user interface screen also includes buttons for selecting one or more agents, requesting a recommendation, requesting graphical analysis of an agent and a button to order the agents according to any desired criteria. In particular, a user might look for the oldest agents in the system, i.e., the most robust agents, or look for agents having certain predictability values, or identify agents that are relatively conservative or aggressive. After identifying and selecting the desired agents, a user preferably uses those agents, thereby tailoring his investment strategy in accordance with his own investing philosophy. Ordering of agents is discussed in more detail later herein in the discussion of agent distribution.

Figure 7 shows the ticker user interface screen. For this particular screen, management system 125 monitors and analyzes the various agents to identify stocks for which there is a pair of agents that have complimentary characteristics. For example, management system 125 searches for one agent having a short horizon and one agent that is more long term, yet both recommending buy (or sell, etc.) and both having good predictability values. If such complimentary agents are found, then management system 125 "opens" a position based on the recommendations of the two agents and lists the instrument or stock in the ticker user interface screen of Figure 7.

In the "open positions" portion of the screen are listed the positions that actually are still open. In other words, these complimentary agents are steadily recommending a particular course of action. Thus, listed is the instrument, the

recommendation, and a Q factor, which is an indicator of the strength of the recommendations. Preferably, three stars indicates the strongest recommendation, two stars medium and one star indicates a fairly weak recommendation. The Q factor preferably takes into account the characteristics of the two complimentary agents that decide to open a position, such as their prediction horizons and PFI, as well as the product of their predictability values. For example three stars preferably are assigned to positions opened by two agents whose multiplied predictability exceeds 0.5, where one agent has a prediction horizon less than 5 trading periods and a positive PFI whilst the other agent has a prediction horizon of 10-20 periods and a negative PFI. A back-testing procedure using historical data is preferably used to determine which type of two-agent consensus should get a one, two or three star rating.

Also listed in Figure 7 is the price at which the position was opened and the date on which it was opened. The agents themselves are then listed so that a user can thereafter look at the CVs of either or both of these agents to learn more about them. Alternatively, the user can use the stock manager or agent manager user interface screens to learn similar information.

The next column shown in Figure 7 is "type." Agent pairs of type A have A brains that are looking for a tendency in a price (generically a linear trend) whereas agent pairs of type B have A brains that are seeking an overbought or oversold technical situation (see above discussion on technical templates). The number appended to the type letter (for example type A-2) indicates how closely the two agents' PFI and prediction horizon fit the specified type. For example, in type A positions it is preferable to have a positive PFI agent with a horizon less than five periods and a negative PFI agent with a horizon of 10-20 periods. If both agents fit this description they will open a type A-2 position. If one of the agents fits the description only approximately the type will be A-1: for example, the horizon might

be 30 periods or the PFI might be neutral (n) instead of negative (-). Type A-0 would indicate that both agents fit the description approximately.

The next column in the ticker user interface screen shown in Figure 7 is “current,” which indicates the current price of the instrument. Thus, a position is opened at a certain price and, thereafter, price changes can be monitored. The last column of this screen is a percent profit based on the price at which the position was opened and the current price.

The “closed positions” window of the ticker user interface screen preferably includes similar data as compared to the “open positions” window, but also includes the “close” date indicating when the position on the particular instrument was closed. As mentioned above, positions are opened when both agents in a complementary pair are in agreement to buy or sell. A position preferably is closed under two different circumstances: i) one of the agents changes opinion, e.g., originally both agents said sell and now one says buy; or ii) the predictability of one or both of the agents falls below a certain threshold.

Figure 8 illustrates the agent round table user interface screen. This screen provides summary-like information with respect to each instrument. In the first column, all of the instruments available in the system are listed. Then, for each instrument there is listed its price, the number of agents active for that stock and the overall predictability of the stock, i.e. the sum of the predictability values of all agents for that instrument. In addition there is a column for “PFI”. As explained above, the price feedback indicator preferably is a measure of what the expected reaction of the market will be to the recommendation of the agent. In other words, a buy could lead to a rise in the price of the bought instrument, which might lead the agent to continue to recommend a buy, which is a potentially unstable situation. A negative price feedback indicator is just the opposite. In this case, the recommendation is likely to result in equilibrium rather than instability.

The next three columns shown in Figure 8 are under a main heading of "head of trading." These three columns indicate the weighted average predictability value for the 5 (or fewer) agents per instrument. Weights in this averaging are determined by maximizing the yield/volatility ratio of a virtual portfolio that is constructed by assigning each agent an amount of money equal to its weight. Accordingly, this is an indication of consensus among the agents that seeks to minimize the volatility of a portfolio that adopts these consensus decisions.

Force is a number that varies between minus one and one, where one is a strong buy and minus one is a strong sell. This force is computed from the agents' recommendations by, for example, using the weighted average of the recommendations where a sell recommendation is represented as minus one, a buy recommendation is plus one and a neutral recommendation is zero, and the weights in this weighted average are the same as those described above in the context of the average predictability.

The head of trading ultimately provides an overall recommendation based on the product of the force and predictability described in the last two paragraphs: when this product is greater than one the recommendation is 'buy', it is 'sell' when the product is less than minus one and neutral between minus one and plus one.

Finally, the agent round table user interface screen provides an overall sum of the predictability values for all agents, thereby giving an indication of the overall predictability for all instruments for which agents are available and thereby indicating the degree of predictability in the market as a whole.

Figure 9 illustrates the stock manager user interface screen including an agent recommendation window in accordance with the present invention. As can be readily seen, the information provided in this window is information that is available from other screens, but is provided here for the convenience of the user. It is noted that the column entitled "force" is not the same force displayed in the agent

roundtable user interface screen. The “force” of the stock manager’s recommendation window preferably is a number that varies between zero and a hundred, where zero is a strong sell recommendation and a hundred is a strong buy recommendation. This force value is obtained from the output of a feedforward neural network, multiplied by a factor of one hundred. Preferably, the input to the neural network is the A-brain preprocessed signal, which is described above with respect to the agent factory.

A force of 50 preferably indicates a neutral recommendation. The recommendation is “buy” if the force exceeds a threshold value that is itself greater than 50, or to “sell” when the force drops below a selling threshold. Both decision thresholds preferably are fixed, when the agent is originally created, to a value that maximizes the agent’s profits on past data.

Figure 10 illustrates exemplary graphical analyses windows for a particular instrument and a particular agent, in this case TELMEX L and TELMEXL-73, respectively. Preferably, five different windows with respective curves are simultaneously displayed. Of course, each curve by itself or other combinations of the curves may also displayed, as desired.

One curve is portfolio position, which shows the actual positions taken by the agent both during its training period in agent factory 115 and during its time with the user. A review of the agent’s CV (Figure 4) identifies the agent’s release or graduation date from the agent factory.

A force curve is also shown. The values the graph takes are the same as those encountered for Force in the agent recommendation box in the stock manager user interface screen. Generally, a force of one hundred indicates a strong buy while zero indicates a strong sell (see above).

The neural activity curve provides a measure of how many previous patterns the agent is recalling when a new price is introduced. That is, as new data is

received, the agent tries to remember from the past how many times a similar pattern has been experienced. The more recognizable a pattern, the greater the measure of “avalanches” of brain activity. It directly measures the number of training patterns the B brain of the agent relearns when presented with a new pattern.

The next graph in the middle right of the graphical analysis user interface screen shown in Figure 10 is a price curve, which displays the price of the instrument over time.

Finally, at the bottom right of Figure 10 a portfolio value curve is shown. This curve indicates how much the agent is earning with its particular strategy. Although not shown, by superimposing the price curve on top of portfolio value curve, for example, the agent’s earnings can be evaluated directly. Further, two parallel lines envelop the portfolio curve. These two lines designate a trading canal (mentioned earlier), which is a rough pictorial way of representing what is meant by predictability. That is, the more an agent’s strategy keeps earnings within this trading canal without extending beyond either the upper or lower boundary, the greater the agent’s predictability value. Thus, at least one criterion for determining whether a particular agent 110 should be returned to agent factory 115 for retraining is whether the agent’s earnings are out of bounds of the trading canal. In the preferred embodiment, the lines of the trading canal are calculated to minimize an error function which is equal to the area contained between the two lines, minus ten times the area above the higher of the two lines that is below the portfolio curve, minus ten times the area below the lower line that is above the portfolio curve. Of course, other calculation techniques may be employed to define and create the trading canal lines.

Agent Distribution

The present invention is designed to assist decision makers with making individual correct decisions despite being faced with a complex, dynamic system such as a stock market. In this regard, agent factory 115 responds to changes in the dynamic system by continually generating new artificial agents 110 having, preferably, predictability values that are greater than a minimum threshold, e.g., a value of 1. The life span of any individual agent 110 depends significantly on the rate of change within the dynamic system. That is, in the case of a stock market that is volatile in an unpredictable way, it is likely that a typical agent's life span will be shorter as compared to a situation wherein a stock market is in a less volatile, predictable state.

In accordance with the present invention, a rental or leasing service is the preferable paradigm for agent distribution. Specifically, an investor desiring access to the agents of the present invention is preferably provided with management system 125, upon entering into a contract or leasing agreement. Management system 125 preferably resides on an electronic computer to which the investor has access either directly or over a private or public (e.g., the Internet) network.

Then, since the agents of the present invention have limited lifetimes in view of the dynamic nature of a stock market, for example, the investor preferably agrees to subscribe for a predetermined period of time (e.g., a year, or a month) to the services provided by the present invention. Specifically, once a fee has been paid and/or a contract signed, the investor becomes eligible to receive artificial agents 110 from agent factory 115. The preferred method of agent distribution is by public network such as the Internet. Those skilled in the art will appreciate that, in view of the public network transmission of the agents, encryption may be used as desired to secure their transmission to management system 125. Of course, agents may also be distributed to management system 125 via conventional storage media such as a CD ROM or even a diskette. The only requirement is that the agents arrive at management system 125 intact and uncorrupted. Agents may be updated with a

frequency appropriate to their method of delivery. For example, if delivery is by CD or diskette once every five working days is a preferred delivery rate. For an online implementation no real extra expense is incurred by more frequent updating and, thus, updating can be more frequent.

As agents are retired due to changing market conditions, agent factory 115 preferably automatically creates new agents, which are made available to the subscribing investor or user.

It is often the case that an investor is not interested in subscribing to agents for all stocks, for example for reasons of cost, that are listed on a stock exchange. The present invention therefore permits limited distribution of agents. Thus, the number and particular agents to which a user can have access can be easily controlled.

Agent Editor Tool

The agent editor is a system and method allowing a user at a remote terminal to create agents that implement a trading strategy desired by a particular user and is described with reference to Figures 11 and 12. A user interface at a plurality of terminals allows users to input parameters 5 that may include the agent's aggressiveness, price feedback indicator, prediction horizon, trading validity period among others or previously selected values. The graphical user interface produces a simple trading strategy 7 that matches these parameters to assist the user in the selection of appropriate parameter values. Once final parameter values are decided upon, an application for agent training is submitted 6 preferably via a web interface to a remote computing facility which accesses a large database of financial time series to train the specifically requested agent to produce statistically optimal trading recommendations. The trained agent is then made available to the remote terminal that submitted the agent-behavior parameters and will provide

trading recommendations and scenarios (Figure 12) for this remote-terminal user. Remote-terminal users can view the historical price patterns that were used as training samples to optimize the agent, and in each case see what happened to the stock's price after that pattern developed further. Statistical characteristics of the training sample-set preferably are displayed (though not shown in Figure 12), such as a 90% confidence price range forecast, mean expected price at the prediction horizon, etc. Agents are ranked according to their predictability and retired when this performance falls below a given threshold. Disabled agents appear in the interface with the mention "retired" and no longer provide new trading recommendations.

The artificial agents of the present invention preferably are self-monitoring whereby the agent itself has the intelligence to determine that its recommendations are no longer deemed to be predictable. Figure 13 illustrates a preferable process by which an agent can be disabled and returned to agent factory 115 for retraining. At step 1300 the agent is presented with past examples of input data in order to determine what recommendations the agent would have made in the past and at step 1310 the data is used to test the validity of the agent's decisions (recommendations). In the preferred embodiment of the invention the agent's past decisions are used to create a portfolio value curve which is the result of applying the agent's recommended trading strategy and analyzing the geometrical properties of this value curve by parameterizing the way this curve fits into a canal constituted of two parallel lines, i.e. the trading canal. Based on this validity testing a predictability value p is inferred using the predictability landscape.

Then, at step 1330, it is determined whether p is below a predetermined threshold, e.g. zero. If not, the agent's recommendation is displayed in the appropriate user interface screen. If the agent's predictability is less than the predetermined threshold, then the agent preferably is disabled at step 1350 and the user preferably is warned that the agent needs retraining, step 1360. This process

may also be implemented in agent factory 115 whereby a user is subsequently notified that a particular agent is no longer available. Management system 125 could also include the foregoing functionality whereby agent factory 115 need not be involved.

In view of all of the foregoing, it is evident that the present invention provides an improved, affordable and timely artificial agent creation, distribution and management method and system, which provides a user with a better basis for decision making.

The foregoing disclosure of embodiments of the present invention and specific examples illustrating the present invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many variations and modifications of the embodiments described herein will be obvious to one of ordinary skill in the art in light of the above disclosure. For example, while the agent factory and management system components have been described as separate components, they could be combined into a single system that resides in one location.

Accordingly, the scope of the invention is to be defined only by the claims appended hereto, and by their equivalents.